Signal separation by nonlinear projections: The fetal electrocardiogram

Thomas Schreiber
Physics Department, University of Wuppertal, D-42097 Wuppertal, Germany

Daniel T. Kaplan
Centre for Nonlinear Dynamics in Physiology and Medicine, McGill University,
3655 Drummond Street, Montréal, Québec H3G 1Y6, Canada
(Received 2 January 1996)

We apply a locally linear projection technique which has been developed for noise reduction in deterministically chaotic signals to extract the fetal component from scalar maternal electrocardiographic (ECG) recordings. Although we do not expect the maternal ECG to be deterministic chaotic, typical signals are effectively confined to a lower-dimensional manifold when embedded in delay space. The method is capable of extracting fetal heart rate even when the fetal component and the noise are of comparable amplitude. If the noise is small, more details of the fetal ECG, like P and T waves, can be recovered. [S1063-651X(96)50405-8]

PACS number(s): 87.10.+e, 05.45.+b

Filtering a time series is the process of transforming a single input into two output time series. In the case of noise reduction, one of the output time series contains the signal of interest and the other is noise. In the case of signal separation, both outputs may themselves be of interest. In this paper, we are interested in a practical signal separation problem: how to extract the electrocardiogram (ECG) of a fetus from the ECG of the pregnant mother carrying the fetus. This problem has medical applications to fetal monitoring.

Each method for filtering time series provides a means of specifying what are the differences between the two outputs. Linear filtering implicitly assumes that the difference can be expressed in terms of the Fourier spectral content of the output signals. The linear technique is extremely general and powerful. It also has the advantage that a more-or-less automatic procedure can be used to design an optimal linear filter — this is the Wiener filter. But despite the generality of the linear filtering approach, it appears quite Procrustean when confronted with a problem where the output signals have similar spectral contents.

Nonlinear dynamics has suggested a new class of filtering techniques that are based on the idea of an attractor. This class of techniques can be dramatically superior to linear filtering methods when performing noise reduction on chaotic systems. The methods can also be successfully applied to nonchaotic systems where the signal can be approximated by dynamics on a low-dimensional attractor. For example, we recently demonstrated that nonlinear dynamics filtering methods are quite powerful on the ECG [1].

In the signal separation problem, it is necessary that one can express the dissimilarity between the two output signals using the ''language'' that is implicitly provided by the filtering technique. For instance, in linear filtering, the difference between the two outputs must be expressed in terms of their different frequency contents. Here, we examine the expressiveness of nonlinear dynamics filtering methods in the context of the problem of separating the ECG of a fetus from that of the pregnant mother.

The ECG is a recording of the electrical potential difference between two points on the body surface. This potential difference arises from the extracellular currents produced from cardiac muscle cells in the course of their activation-contraction-relaxation cycle. A pregnant woman's electrocardiogram reflects the activity of the hearts of both the mother and the fetus. (See Fig 1.) Since the fetal heart is much smaller than the maternal heart, the fetal signal is typically much smaller than the maternal. Normally, the fetal heart beats faster than the maternal heart, and the times of maternal and fetal beats are considered to be independent of one another. The fetal ECG is otherwise quite similar to the maternal ECG, and in particular both have broadband spectra.

By positioning the ECG electrodes on the abdomen, one can maximize the amplitude of the fetal component of the ECG, but it will still be difficult to extract the fetal signal,

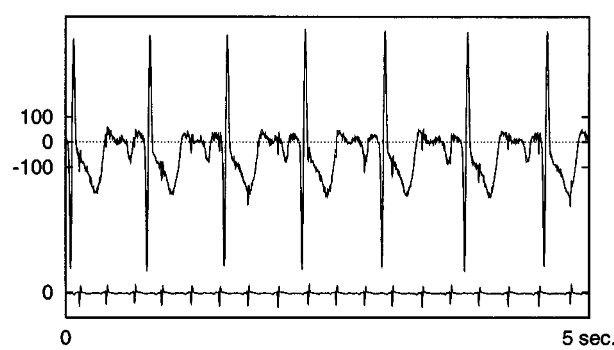

FIG. 1. A synthesized signal consisting of the sum of two actual ECGs from the MIT-BIH Polysomnographic Database [3]. One of the ECGs has been sped up by a factor of 3 and reduced to 0.07 in amplitude in order to mimic a fetal ECG. In addition, the series has been contaminated by typical base line noise of the same rms amplitude as that of the fetal component. We also show the pure fetal ECG included in the signal. Note that the figure only shows one quarter of the time series we used for the signal separation procedure. The sampling rate is 250 Hz. Due to the 12 bit A-D conversion, the units here are not simply μV.

1063-651X/96/53(5)/4326(4)/$10.00 53 R4326

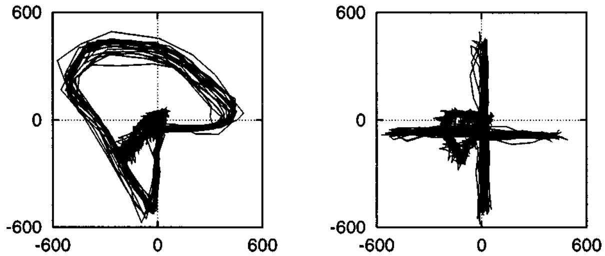

FIG. 2. Delay representations of the ECG shown in Fig. 1. Left: The delay time is 0.02 sec. and the large loop resolves the ventricular QRS complex while the smaller features in the center contain the atrial P wave and the T wave. Right: When a longer delay time (0.08 sec.) is used, the QRS complex is no longer resolved but the small central loop displays the T wave. (See [7] for an introduction to the terminology and electrocardiography in general.) In delay coordinates, the fetal component is hardly visible as small deviations from the manifold.

FIG. 3. A synthesized maternal-fetal ECG for an almost noise free situation where however the fetal component is very small. Both ECGs in this figure were taken by Petr Saparin at Saratov State University.

especially during the spikelike QRS complex. A number of methods have been proposed to perform this extraction other than by visual inspection. Most approaches use multiple leads which are subtracted in an appropriate way in order to approximately cancel out the maternal signal. If a single lead is used, one can subtract a template of an averaged ECG cycle from the maternal signal [2]. Filtering in the Fourier domain is not appropriate since all three components, maternal, fetal, and noise, have broad spectra that are hard to separate.

The strategy pursued in this paper is based on an observation which has been discussed and used for ECG noise reduction in previous work [1]: Although not deterministically chaotic, the ECG can be unfolded using delay coordinates and is then empirically found to be effectively constrained to a low-dimensional manifold. This is related to the relatively good predictability of the ECG over times smaller than a cycle length. Such a delay representation is shown in Fig. 2. (See [4] about the theory of time series embeddings.) In [1], noise reduction has been accomplished by constructing locally linear approximations to this manifold and projecting onto these linear subspaces. In the present paper, the fetal component will be treated similarly to the noise, that is, as a disturbance of the maternal signal.

We will again use the local projective noise reduction algorithm proposed in [5] and adapted to the ECG case in [1]. (See [6] for a review on nonlinear noise reduction.) For a technical description of the algorithm we refer to these references. Let us, however, recall that the algorithm requires the choice of three important parameters: the length of the embedding window, the dimension of the local manifold projected on, and the diameter of the neighborhoods used to form the linear approximations. The embedding window can be used to select components by time scale, the neighborhood size determines a length scale in phase space.

These parameters constitute the ''language'' that allows us to express the differences between the fetal ECG and the maternal ECG. In our case, the relevant medical facts are that the fetal ECG is much smaller in amplitude than the maternal ECG, and that the fetal heart beat occurs on a shorter time scale than the maternal one. We set the nonlinear filtering parameters accordingly: the embedding window is set to 0.2 sec. corresponding to a small fraction of a maternal heart beat cycle but covering practically an entire fetal beat; the local neighborhoods are set large enough to contain the peak-to-peak amplitude of the fetal ECG. As a result, the locally linear projections are inadequate to preserve the structure of the fetal ECG: essentially the whole complicated nonlinear structure of a fetal beat is being projected onto a linear manifold. The output of the nonlinear filter is therefore stripped of the fetal ECG component. To recover the fetal ECG, then, one can subtract the output of the filter from the input, as shown in Figs. 3 and 4. In such a case, although the fetal component is quite small, it can be recovered from the superimposed series in such detail that the estimation of clinically relevant parameters, like the P-Q interval, is possible.

The idea of discriminating between two signals based on their peak-to-peak amplitudes is quite foreign to the linear filtering methodology, which has no effective means of expressing this simple difference between the fetal and mater-

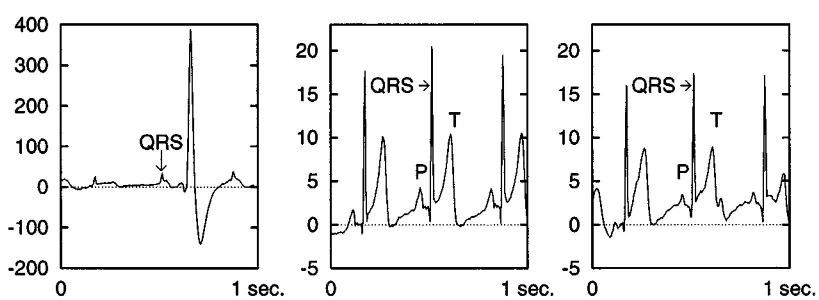

FIG. 4. Detail of the input signal (left), the fetal component (middle), and the reconstructed fetal component (right). Same series as in Fig. 3. We also identified some clinically relevant features of the fetal part: the P wave indicates the depolarization of the atrium. The QRS complex reflects the depolarization and the T wave the repolarization of the ventricle.

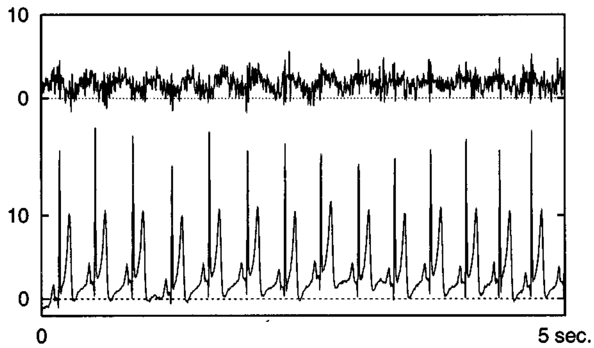

FIG. 5. The result of Wiener filtering (top) is shown in comparison with the true fetal component (bottom). Same data as in Figs. 3 and 4. Although we used the known spectrum of the fetal component to construct the filter, the optimal linear filter is essentially useless for extracting the fetal signal. The large peaks in the Wiener filter output correspond to the maternal QRS complexes.

nal ECG signals. The use of amplitude as a discriminating feature also leads to a paradox: the performance of the filter can become worse as the amplitude of the fetal signal increases.

We can compare the nonlinear dynamics filtering process to the optimal linear filter, the Wiener filter. To construct the optimal linear filter, we need estimates of the power spectra of the maternal and fetal signals separately. For the purposes of demonstration, we will use the actual fetal and maternal signals to estimate the power spectrum, recognizing that the resulting filter will be more refined and precise than would be possible in practice. The result of the Wiener filter thus constructed is shown in Fig. 5.

The signal shown in Fig. 4 is unusually noise free. In a more typical case, there is considerable content in the ECG that reflects neither the maternal heart nor the fetal heart, and that we consider exogenous noise. In addition, there may be physiological variability in the maternal ECG that cannot be well represented by dynamics on a low-dimensional manifold and that therefore shows up in the extracted fetal signal. We wish to try to estimate this physiological variability and the noise in order to minimize the contamination of the extracted fetal signal. This can be done by performing noise reduction with a shorter embedding window and a smaller neighborhood size. The objective now is to preserve the fetal signal which is possible because using the shorter embedding window and smaller neighborhoods allows the locally linear projections better to represent the fetal components. The difference between the input and the output is now taken to indicate the component of the ECG which is neither of fetal nor maternal origin. As seen in Fig. 7, middle strip, this approach is only partially effective in practice: although the maternal and fetal components are greatly reduced in amplitude, there are remnants particularly during the fetal and maternal QRS complexes.

Combining the outputs of both filters, as indicated in Fig. 6, appears to give the best results. (See Fig. 7.) Although some part of the fetal signal is lost, the combination compensates for the inability of the large-neighborhood filtering perfectly to represent the peaks of the maternal QRS complex. In the case of the noisy ECG, only the QRS complex of the fetal ECG is faithfully extracted by the nonlinear filter. Details such as the P wave are lost. In a clinical situation, fetal monitoring is mainly concerned with fetal heart rate, and so the timing of the QRS complex is the important feature that must be preserved in the signal separation process. This clinical criterion is satisfied even in the noisy ECG case.

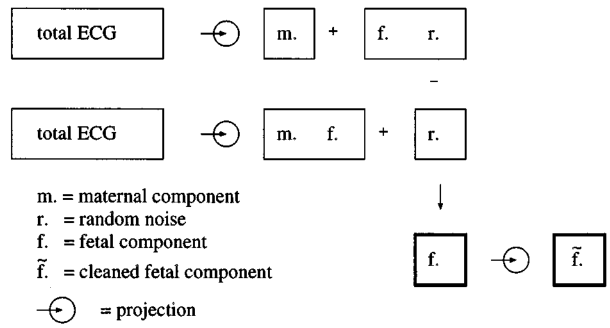

FIG. 6. Schematic diagram for the extraction of the fetal ECG. Two locally linear projections are performed. The fetal component is estimated as the difference of the two resulting signals (or ''noises''). Sometimes it can be useful to perform an additional noise reduction step on the estimated fetal ECG.

Technically, the first filtering step is done with an embedding window that is somewhat shorter than one ECG cycle, 0.2 sec. in the examples in this paper. Further, the diameter of the neighborhoods is chosen large enough in order to cover both the fetal signal and the noise. For the signal in Fig. 1 we used neighborhoods of diameter analog-to-digital (A-D) 50 units, while for the smaller fetal signal in Fig. 3 we chose 25 A-D units. This relevant length scale was determined by inspection of the measured signal. The difference between the input and the output of this filter contains fetal and noise disturbances of the maternal ECG.

In order to select the exogenous noise, the embedding window of the second filter is much shorter (here 0.02 sec.) such that the fetal component is interpreted as part of the

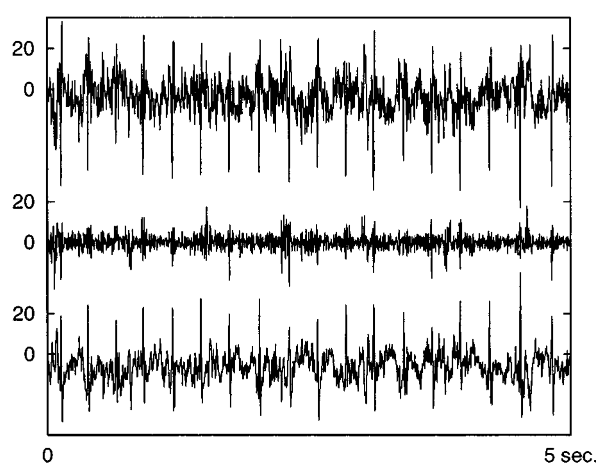

FIG. 7. Top: all contaminations of the maternal signal, Fig. 1, containing fetus and noise. Middle: attempt to select the random contamination only. Bottom: difference of the two previous panels.

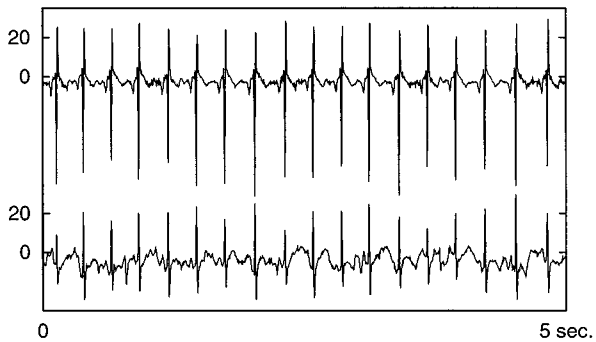

FIG. 8. Top: original fetal component included in the series shown in Fig. 1. Bottom: reconstructed fetal component after nonlinear noise reduction. Although the amplitude of the fetal *QRS* complex is reduced in the reconstruction, at least the heart rate can be determined reliably.

signal. The neighborhoods are smaller (10 A-D units in Fig. 1) as well in order to preserve the fetal contribution. The difference signal, Fig. 7, bottom, was then further processed using a delay window of 0.1 sec. and neighborhoods of diameter 10 A-D units. The final version of the reconstructed fetal ECG is shown in Fig. 8. We are not particularly interested in the maternal signal itself. Thus, since we form the difference of two reconstructions, some smooth overall distortion of the maternal signal is of no concern, and we always project onto a one-dimensional manifold, except for the last processing of the fetal signal, where projections onto three dimensions were used. We used a delay of 4 msec for all embeddings. This usually gives the best filter properties. The necessary embedding dimension and thus the computational effort can, however, be reduced dramatically when a longer delay time is chosen. In that case, the filtering can be done in about real time on an Alpha 200 workstation, once a good set of parameters is established.

A number of different approaches have been developed for extracting the fetal ECG [2]. These approaches are all nonlinear, but generally consist of *ad hoc* techniques for identifying the maternal QRS complexes and subtracting out templates reflecting the typical maternal QRS complex. For instance, a template might be constructed by averaging together many aligned maternal complete P-QRS-T cycles. The approach we present here offers a natural framework for this sort of subtraction: rather than using a template in the time domain, we estimate one in the domain of the embedded signal. The template takes the form of a low-dimensional manifold. This domain allows us to express the differences between the fetal and maternal components in terms of amplitude and time scale, which are closely based on what we know about the physical differences between the fetal and maternal ECG.

We thank Petr Saparin for providing ECG recordings and Peter Grassberger and Holger Kantz for stimulating discussions. This work was supported by the SFB 237 of the Deutsche Forschungsgemeinschaft, the Natural Sciences and Engineering Research Council of Canada, and the Fonds de la recherche en santé du Québec.

[1] T. Schreiber and D. T. Kaplan, CHAOS **6**, 87 (1996).

[2] P. P. Kanjilal, S. Polit, and P. K. Dey, in *Proceedings of the 12th Triennial World Congress, International Federation of Automatic Control, Vol. 3, Sydney NSW, Australia 1993* (Pergamon, Oxford, U.K., 1993), pp. 183–186; E. Cicinelli, A. Bartone, I. Carbonara, G. Incampo, M. Bachicchio, G. Ventura, S. Montanaro, and G. Aloisio, Int. J. Bio-Med. Comput. **35**, 193 (1994); J. R. Mazzeo, Med. Prog. through Technol. **20**, 75 (1994); Y.C. Park, K.Y. Lee, D.H. Youn, N.H. Kim, W.K. Kim, and S.H. Park, IEEE Trans. Biomed. Eng. **39**, 868 (1992); S. Abboud, A. Alaluf, S. Einav, and D. Sadeh, Comp. Biol. Med. **22**, 325 (1995); Y. Tal and S. Akselrod, Comp. Biomed. Res. **24**, 296 (1991); S. Abboud, G. Barkai, S. Mashiach, and D. Sadeh, Comp. Biomed. Res. **20**, 147 (1990).

[3] MIT–Beth-Israel Hospital Polysomnographic Database, Biomedical Engineering Centre, MIT, Cambridge (1991).

[4] T. Sauer, J. Yorke, and M. Casdagli, J. Stat. Phys. **65**, 579 (1991).

[5] P. Grassberger, R. Hegger, H. Kantz, C. Schaffrath, and T. Schreiber, CHAOS **3**, 127 (1993).

[6] E. J. Kostelich and T. Schreiber, Phys. Rev. E **48**, 1752 (1993).

[7] A. L. Goldberger and E. Goldberger, *Clinical Electrocardiography* (Mosby, St. Louis, 1977).